DOCUMENT RESUME

ED 454 863 IR 058 154

AUTHOR	Moen, William E.
TITLE	Resource Discovery Using Z39.50: Promise and Reality.
PUB DATE	2000-11-00
NOTE	21p.; In: Bicentennial Conference on Bibliographic Control for the New Millennium: Confronting the Challenges of Networked Resources and the Web (Washington, DC, November 15-17, 2000); see IR 058 144.
AVAILABLE FROM	For full text: http://lcweb.loc.gov/catdir/bibcontrol/moen_paper.html.
PUB TYPE	Reports - Descriptive (141) -- Speeches/Meeting Papers (150)
EDRS PRICE	MF01/PC01 Plus Postage.
DESCRIPTORS	*Information Networks; *Information Retrieval; *Information Systems; Models; Semantics; *Standards
IDENTIFIERS	Interoperability; Search Engines; *Z39 50 Protocol

ABSTRACT

This paper provides a portrayal of Z39.50 (a communications protocol that enables two systems to exchange messages for the purpose of information retrieval) that explains its flexibility in response to a variety of information retrieval requirements in the networked environment. The first section presents an overview of Z39.50, including the evolving context for the protocol, the abstract model for information retrieval, the Z39.50 model, and semantics for searching and retrieval. The second section discusses resource discovery, including a comparison of search engines and Z39.50, granularity and aggregation, and the Z39.50 model of resource discovery. The third section addresses interoperability, including levels and types of interoperability, semantics for searching, and Z39.50 profiles as solutions to semantic interoperability. The fourth section covers the application of Z39.50 in virtual union catalogs and cross-domain searching. The final section considers Z39.50's future in networked information retrieval. (Contains 23 references.) (MES)

Reproductions supplied by EDRS are the best that can be made
from the original document.

Resource Discovery Using Z39.50: Promise and Reality

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)
This document has been reproduced as received from the person or organization originating it.
Minor changes have been made to improve reproduction quality.
Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY
B. Wiggins
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

ED 454 863

William E. Moen
Assistant Professor
School of Library and Information Sciences
University of North Texas
Denton, Texas 76203

Final version

The ANSI/NISO Z39.50 protocol for information retrieval addresses the complex challenges of intersystem communication. Original uses envisioned for the protocol look very little like current implementations and uses. In the 1980s, users on one library catalog system would search and retrieve bibliographic records on a remote system. By the late 1990s, there was a need for discovering networked resources and integrating access to them. Yet, the Z39.50 protocol has addressed both these scenarios. This paper provides a portrayal of Z39.50 that explains its flexibility in response to a variety of information retrieval requirements in the networked environment.

What Is Z39.50 Really?

At its most basic, Z39.50 is a communications protocol that enables two systems to exchange messages for the purpose information retrieval. However, one can define and characterize Z39.50 in a number of ways. To begin to understand the use of Z39.50 today, it is worth a brief look back over its 20+ year history [1]. Z39.50 was a realization of 1970s visions for connecting computer systems of large bibliographic utilities and research libraries via telecommunications for purposes of resource sharing, specifically, for sharing MARC bibliographic and authority records. Library leaders such as Henriette Avram saw the potential for resource sharing through the convergence of telecommunications and computers, thus moving towards a regime of national bibliographic control. The National Information Standards Organization (NISO) [2] established Subcommittee D in 1979 to develop a "computer–to–computer protocol for electronic communication of digital information over a network" to support "information transfer at the application level" and would depend on other standards for underlying protocol layers [3]. The Subcommittee focused its initial effort on a protocol for information retrieval.

An Evolving Context for the Protocol

Technical standards can be viewed as solutions to problems. In the case of Z39.50, one can ask what problem was being addressed by the information retrieval protocol. Libraries were the context for the problem. The problem was how to get diverse library automation systems and their underlying information retrieval systems to communicate and thus enable users of one system to search another library's catalog and retrieve MARC records. In its origins, the protocol was intended to solve library problems.

Through the 1980s as the standards committee continued its work, the centrality of the library problem for intersystem

IR058154

BEST COPY AVAILABLE

communication remained paramount, but new voices became stronger in response to the emerging information retrieval protocol. These voices (e.g., from the abstracting and indexing services) called for a more generalized information retrieval protocol, not one focused only on the intersystem communication between libraries' bibliographic record systems.

With the approval of Z39.50 Version 3 in 1995, the range of implementors of and applications for Z39.50 broadened to include communities with requirements for information retrieval among diverse and distributed resources. Government information, geospatial information, and museum information were three application areas adopting and adapting Z39.50 to the needs of their communities. No longer was the library catalog the central application area for Z39.50.

So, what is Z39.50 really? It is a computer-to-computer protocol that enables intersystem communication for the purpose of searching and retrieving information (where the information can be in the form of MARC records, data from geospatial datasets, museum object records, etc.). But that does not explain why a standard that developed in the context of library problems is now used in a variety of other communities and their applications. For that, we need to look at what the standard offers.

Models, Semantics, and Bits on the Wire

Anyone picking up the Z39.50 standard with the goal of learning what it is, what it does, and how it does it is usually disappointed. Instead of clear descriptions of Z39.50's capabilities and practical uses, the reader is confronted by complex and abstract technical descriptions of facilities, services, application protocol data units, parameters, option bits, and ASN.1 structures. Without initiation into this technical language, the document remains opaque. Yet that technical language does more than confound the average reader. It expresses three important components that are central to what Z39.50 is:

- Abstract models of information retrieval activities (e.g., search, retrieval, etc.)
- A language consisting of syntax and semantics for information retrieval that enables communication between systems
- A prescription for encoding search queries and retrieval results for transmission over a network infrastructure.

Focusing on these components allows us to see the strengths and limitations of Z39.50 for networked information retrieval.

A major contribution of the standard is an abstract model of information retrieval [4]. As an abstract model, it is not tied to any specific implementation, database design, or search engine. Wake states that the "complexity of the Z39.50 information retrieval model should be seen as richness that enables this model to describe many retrieval systems" [5]. The components of the model include (see Figure 1):

- Query: the search submitted by the user (for details about the query, see below on semantics) from a client
- Database: the physical or logical repository of records
- Database record: a local data structure within a database
- Result set: a list created by the server of pointers to database records that meet the criteria of the query
- Retrieval record: the data from the local database record formatted for interchange in a syntax understood by both systems.

This model allowed Z39.50 protocol developers to conceptually separate the user interface (for formulating searches and displaying results) from the information server (with its database management system, search engine and algorithms, local record structure, etc.). Z39.50 protocol machinery in the form of Z39.50 clients and servers mediates between two systems as represented in Figure 2. But for this model to be effective in intersystem communication, protocol developers needed to agree on a language that Z39.50 clients and servers would *speak* to carry out information retrieval transactions.

Figure 1
Abstract Model of Information Retrieval

Abstract Model of Information Retrieval

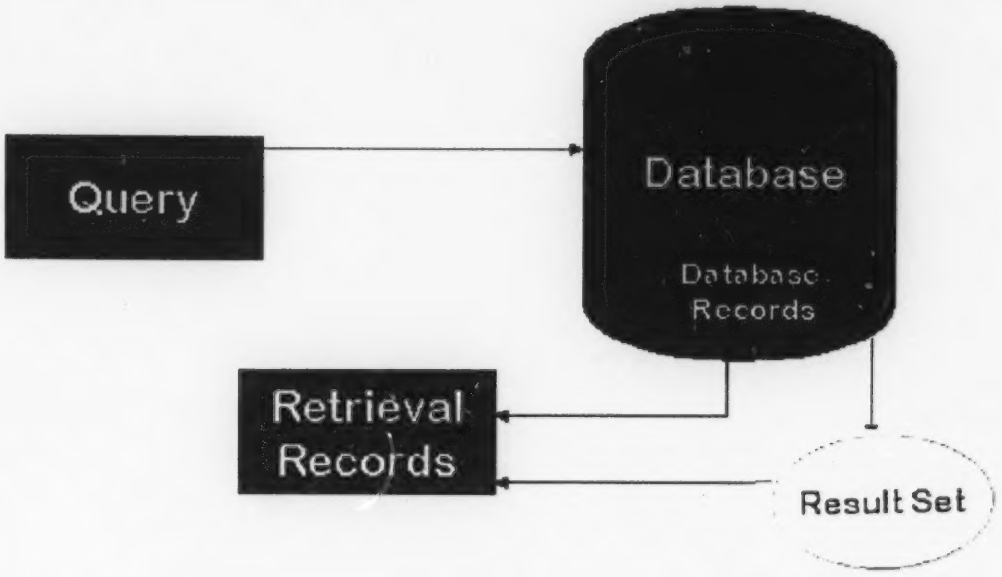

Figure 2
Z39.50 Model of Information Retrieval

Z39.50 Model of Information Retrieval

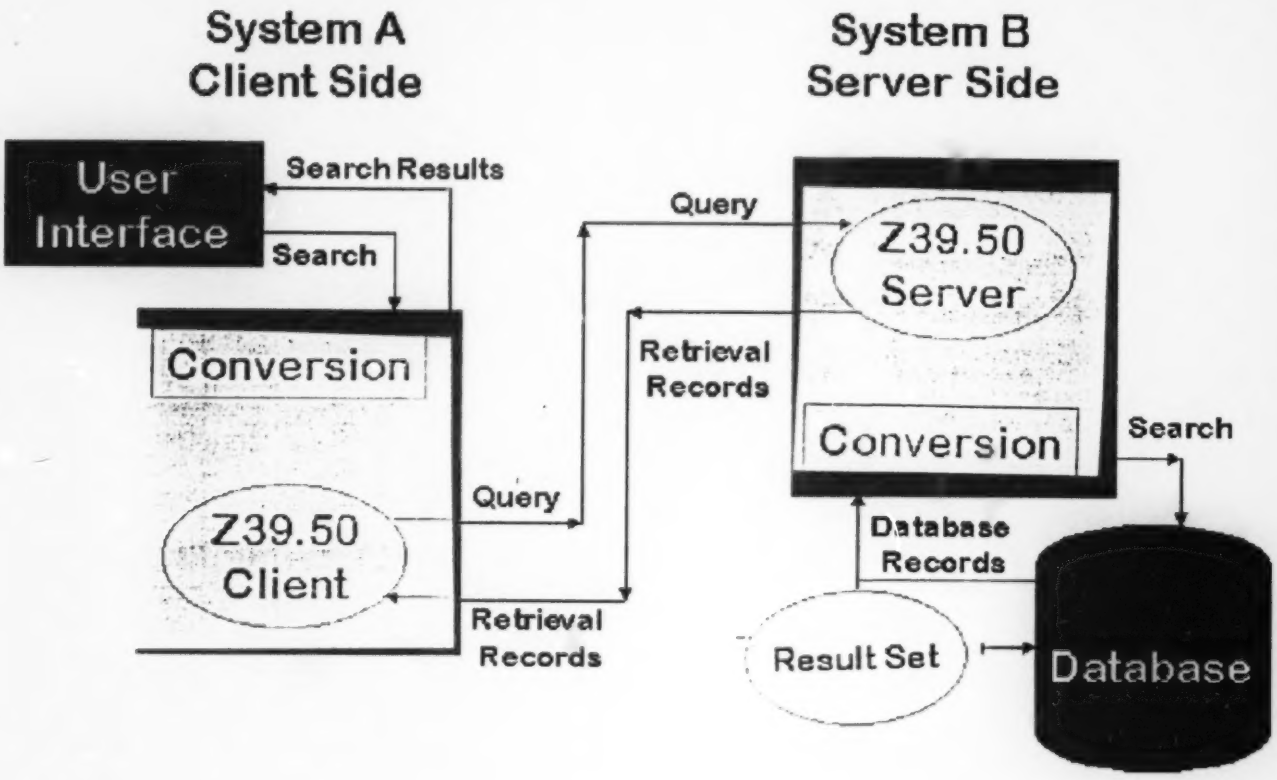

Semantics for Searching and Retrieval

How does a user instruct one system to ask a remote system to do a search for books **by** Mark Twain? How does the remote system know that the query it receives is requesting a search for books **by** Mark Twain and not books **about** Mark Twain. What about a title exact match search? What does a title search mean anyway? These questions point to the second major contribution of Z39.50 developers: a semantic model for expressing searches and requesting records that match the criteria of the searches, and the semantics for interchanging the retrieval records.

Each online catalog with its underlying information retrieval system provides users with various search and retrieval options. Typically, search and retrieval options differ between vendors' products. Achieving communication between these disparate systems, each with their own search and retrieval capabilities, was the challenge faced by Z39.50 developers. Getting two systems to exchange protocol messages is one technical challenge, but getting them to "understand" what the messages mean is the arena of semantic interoperability [6].

BEST COPY AVAILABLE

Building on the abstract model in Figure 1, the developers first worked on standard semantics for expressing queries. More recently, Z39.50 developers focused on semantics and structures for retrieval in a networked information world no longer populated with MARC records. We focus here on semantics for searching to illustrate how Z39.50 addresses semantic interoperability.

In an online catalog environment, users interact with the information retrieval system through an interface where they first formulate their search into a query understood by the machine. A query typically has a search term that is characterized by qualifiers. For example, a search for books by Mark Twain is formulated into a query where the search term is "Mark Twain" (or possibly "Twain, Mark"), and this term is characterized as an "author" term (i.e., search the access point "author"). The qualifiers for the search term tell the information retrieval system how to execute the search: do a search for all records in your database where author is equal to "Mark Twain." We can more precisely characterize the search term and how we want the query executed by additionally describing:

- the structure of the search term (is it a word, a phrase, a date, etc.)
- whether truncation should be performed and if so why kind (no truncation, right truncation, left truncation, etc.)
- whether the search term match the entire field value or only part of the field.

To generalize based on this understanding of what queries are and do, Z39.50 provides attributes sets for expressing searches. Attribute sets define the types of qualifiers available for a search term, and define specific values for those attribute types. For example, the Bib-1 Attribute Set is widely used to express Z39.50 queries against library catalogs. It defines six Attribute Types, each designated by a name and integer: Use(1), Relation(2), Position(3), Structure(4), Truncation(5), and Completeness(6). Each attribute type can take on values (also designated by name and integer value). For example, a Use attribute characterizes the access point that should be searched. One Use attribute value is "Title" or "4" to designate a title access point. Attribute types and values are expressed as integer pairs; the pair (1,4) tells the server to execute a title search. The combination of attribute types and values provides a way to express the semantic intention of the search and prescribe the behavior expected when the server executes the query. For example, we can express a ***keyword author search*** for Twain as ***(1,1003) (2,3) (3,3) (4,2) (5,100) (6,1) Twain***, where:

- Use Attribute (1) = author (1003)
- Relation Attribute (2) = equal (3)
- Position Attribute (3) = any position in field (3)
- Structure Attribute (4) = word (2)
- Truncation Attribute (5) = do not truncate (100)
- Completeness Attribute (6) = incomplete subfield (1).

I've illustrated in some detail how Z39.50 addresses semantic interoperability for searching by providing a standardized language (syntax and semantics) for expressing queries. For meaningful communication to occur, the communicating Z39.50 client and server must "know" or recognize values from a common attribute set (e.g., Bib-1). Only then will they be able to meaningful exchange and process a query. For example, the client will be able to convert a search expressed in the structure of its local information retrieval (IR) system into standard Z39.50 vocabulary; and the server will be able to receive and understand the Z39.50 query and convert it into its local IR system search logic for execution. Figure 2 indicates the conversion points for mapping into and out of the Z39.50 protocol language on the client and server.

The expressiveness offered in Z39.50 for queries grew out of the context for the protocol, namely, searching large online catalogs and bibliographic databases accessible by robust information retrieval systems. These databases held well-structured bibliographic records created according to national and international standards and guidelines. The information retrieval systems provided any number of access points to the records including author, title, and subject, and allowed the end-user to qualify and refine searches to improve retrieval results. The model for searching was not simple keyword access. Z39.50 functionality mirrors the search and retrieval functionality of those online library catalog systems. One power of Z39.50 is being able to communicate precision-oriented (as well as recall-oriented) searches against well-structured information in the form of bibliographic records or other forms of structured metadata. What are the implications of this for resource discovery?

Resource Discovery

We know that resource discovery must be a good thing since lots of people want to do it and many claim they have tools to do it. Like the term metadata, resource discovery has many connotations. To evaluate the use of Z39.50 for resource discovery, it is helpful to have a working definition of the concept. Lynch suggests that the resource discovery is used to describe a complex collection of activities, from "simply locating a well-specified digital object on the network all the way through lengthy iterative research activities....Discovery often involves the searching of various types of directories, catalogs, or other descriptive databases....Most often, the discovery process operates on surrogates (such as descriptions) of actual networked information resources" [7]. Key elements of resource discovery appear to be finding, identifying, and accessing information, and the use of representations or surrogates in the discovery process.

Lynch characterizes networked information resources as "digital objects, collections of digital objects, or information services on the network" [7]. One can use the Internet to discover all kinds of resources, such as people, organizations and institutions, products, services, texts, images, sounds, and so on. Each of these resources are represented digitally in some fashion. People could be represented by the occurrence of their name on a document, in an email message, or on a website. Organizations and institutions might be represented by a company website How these objects are represented will likely determine the utility of Z39.50 for discovering them.

From the perspective of the Z39.50 abstract information retrieval model, there is a database that contains records, where a records is a surrogate for some thing (e.g., a digital object). With Z39.50, a Z39.50 client knows of the existence of a Z39.50 server (e.g., network address, port number, etc.) and possibly names of one or more databases made accessible via the server. This means that to get started with resource discovery using Z39.50, a client must know at least one server. But that is really no different than needing to know the URL for AltaVista or Google to get started doing resource discovery using Web search engines.

Apples and Oranges, Search Engines and Z39.50

One can hardly discuss networked information discovery and Z39.50 without a brief discussion of web search engines. Although it is critical in evaluating Z39.50 role in resource discovery to clarify the differences between Z39.50 and web search engines, the scope of this paper does not allow an extended treatment. Z39.50 is an intersystem communications protocol for information retrieval. It is not a search engine. A Z39.50 client can send searches to one or more database on remote systems at the same time (from the perspective of the user). It allows the user to see these different databases as if they were one logical resource. The client connects with each separate server, searching the current contents of the database, and getting results directly from the source databases. Z39.50 simply provides the protocol for these systems to communicate information retrieval messages. One can characterize this approach to networked information retrieval as decentralized or multi-system.

A web search engine is fundamentally a single information retrieval system that has the added function of harvesting resources from the Internet and performing some sort of indexing to make those resource searchable. When users are interested in discovering resources via a web search engine, their web browser presents a search interface for that search engine, and a query is executed against the databases and indexes of that single search engine. One can characterize this approach to networked information retrieval as centralized or single-system.

The stored representations may differ significantly between a Z39.50 accessible database and the web search engine databases. In the latter, the harvested networked information resources are typically represented by words/terms taken from the document and placed in a index. There is no structured representation for the resources. Z39.50 accessible databases typically contain structured representations or surrogates for the resources. These may be in the form of library catalog bibliographic records, museum object records, collection-level records, or other forms of structured metadata.

Granularity and Aggregation: What are Users Trying to Discover?

We noted above that a Z39.50 client must "know" about a Z39.50 server prior to getting started. There are published lists of Z39.50 servers, but the larger challenge is selecting an appropriate server for a particular information need. Subject gateways, such as the Arts and Humanities Data Service [8], assist users by identifying a number of resources (i.e., databases) that are Z39.50 accessible and provide a Web search interface for using Z39.50 to search one or more of the identified resources at the same time. The gateway is a logical aggregation of several discrete networked information resources. This raises the question as to what the resources are that discovery tools are helping users discover? Web search engines work at the level of an HTML file (the addressable unit for retrieval), where the file can be a report, a homepage, a poem. Z39.50 models resources as records (the addressable unit for retrieval in a database), where the record can represent almost anything that can be described.

The library cataloger's concept of unit of analysis (or unit of description or unit of retrieval) is useful in this context. This concept helps catalogers identify what exactly they are representing in a single bibliographic record. In the print world, various levels of granularity or aggregation can be represented. For example, a single volume of an monographic set can be described in a bibliographic record; the monographic set also can be described.

In terms of resource discovery, what exactly is the *size* or *scope* of the resource we are trying to discover? Are we looking for a web page? A web site? A text document comprising a number of web pages? A specific graphic image that is part of a web page? A database of records? The unit of analysis for web search engines is an addressable file. The unit of analysis for Z39.50 can be anything, but the record-based model for Z39.50 assumes that a resource is represented by a logical, if not a physical, record. Some examples can illustrate this. In a library catalog, a record might represent an item in a library's collection such as a book, journal, map, etc. But a record might also represent a series, a set of items. In an abstracting and indexing service (A&I) database, a record might represent a journal article. In a museum collection management system, a record might represent a specific art object. We can categorize all of these as metadata records, structured records that describe resources. We can also envision descriptive metadata records created to represent an online database, a repository of electronic texts, a museum and collections housed by that museum. This moves us to a context in which Z39.50 can be viewed as a tool for resource discovery. As long as the resources are represented and made available through some sort of information retrieval system, those resources could be discovered via Z39.50. Figure 3 illustrates a Z39.50 client accessing one or more Z39.50 accessible information retrieval systems that have records representing information resources. Z39.50 discovers those resource descriptions. Whether or not the described resources are accessible via Z39.50 or any network tool is another issue.

To accomplish Z39.50 resource discovery, the system represented by the User Interface in Figure 3 must be interoperable with one or more remote information retrieval systems and the databases served by those information retrieval systems so meaningful communication occurs. The challenge is, can a user formulate a search using a Z39.50 client to search one or more remote systems and get meaningful results? This is the fundamental challenge of interoperability.

Figure 3
Z39.50 Model of Resource Discovery

Z39.50 Model of Resource Discovery

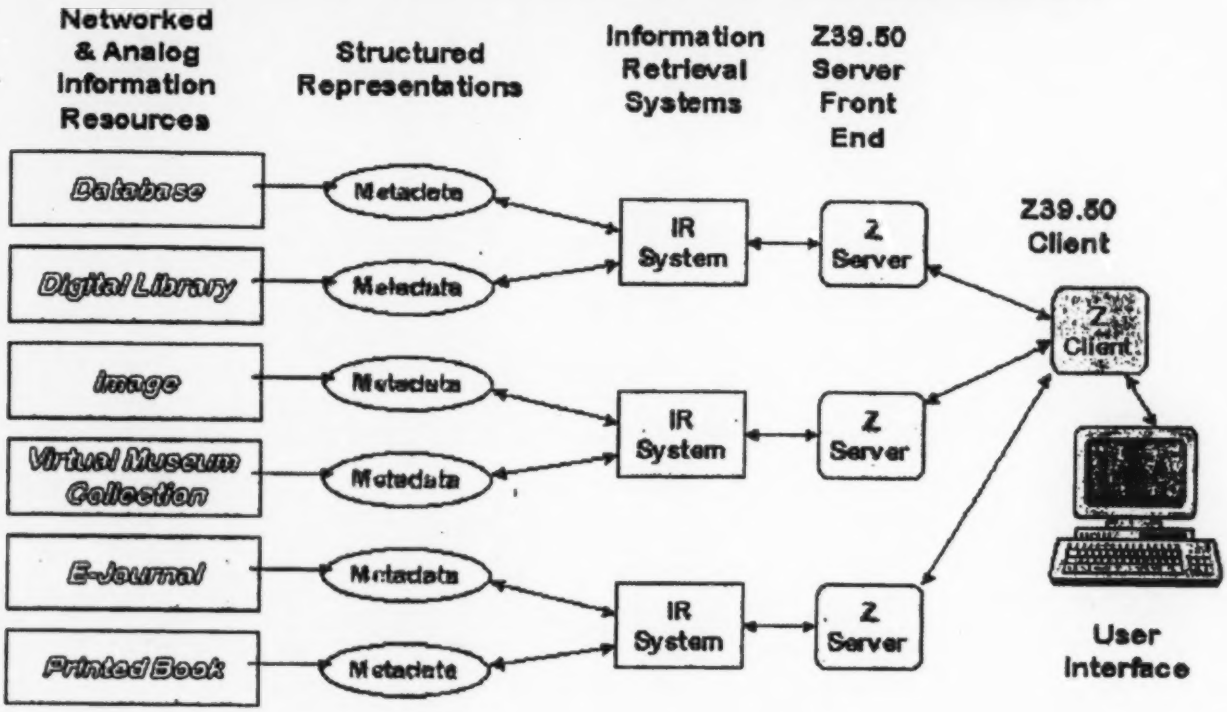

Interoperability

Interoperability is a key issue for resource discovery and more generally networked information retrieval [9]. Interoperability is a concept that addresses the extent to which different types of computers, networks, operating systems, and applications work together effectively to exchange information in a useful and meaningful manner. The networked environment is heterogeneous; it hosts many different technologies, various data, multiple applications, and other networked life–forms. A functional goal in this environment is to hide this heterogeneity from users so they may effectively do business, search for information, communicate, and perform other tasks. There is little doubt interoperability is a key issue in the networked environment [6, 10, 11, 12]. Interoperability or its absence can affect information access. Technical interoperability can raise important policy and organizational issues [13].

BEST COPY AVAILABLE

As a working definition of interoperability for this paper is: *the ability of different types of computers, networks, operating systems, and applications to work together effectively, without prior communication, in order to exchange information in a useful and meaningful manner* [14]. Based on experiences with Z39.50 implementations, several levels and types of interoperability can be articulated including:

- **Low–level protocol (syntactic):** do two implementations interchange protocol messages according to the standard?
- **High–level protocol (functional):** do two implementations support the same Z39.50 services as defined in the standard?
- **Semantic level:** do two implementations preserve and act on meaning of information retrieval tasks?

Z39.50 implementation experience gained over the past decade has solved most of the low-level protocol interoperability problems. The high-level protocol interoperability problems are resolved for the most part when a Z39.50 client and Z39.50 server support the same services (e.g., sort, scan). The arena of semantic interoperability is where Z39.50 developers and implementors face the most complex set of challenges.

Semantics for Searching Revisited

We discussed above how Z39.50 provides a language for expressing queries, and this language with its attendant syntax and semantics, enables two systems to *understand* each others requests and responses. In practice this *understanding* has not always been achieved. The lack of semantic interoperability has caused users to lose confidence in Z39.50 interfaces to information retrieval systems (whether their native systems or remote systems). What affects semantic interoperability? The two major factors affecting interoperability are differences in Z39.50 implementations and differences in indexing decisions in the information retrieval systems. The results of these differences show up in retrieval results. Going back to the analogy of Z39.50 as a language, the meaning (semantics) of the protocol messages needs to be clear if two systems are to share an "understanding" of the message. Z39.50 provides standardized "vocabularies" to express queries using registered sets of attributes (where attributes are used in the Z39.50 query to characterize a search term). The attribute sets provide the "words" in the vocabulary for searching.

Z39.50 implementations, however, do not always support (i.e., understand and act on) the same "words" from the standardized vocabulary for searching. Taking an example from library catalogs, System A wants to search System B for a corporate author and formulates the query using the correct Z39.50 attribute type/value pair to characterize its search term as a corporate author. But System B does not support that particular Z39.50 attribute type/value pair. The semantic intention of the user and his/her search cannot be acted upon. However, the System B does support a name search, and in an attempt to be helpful, processes the corporate author search as a name search; the results, however, may include records that are not relevant to the original corporate author search; semantic loss has occurred. In both these cases, semantic interoperability is reduced or does not exist.

The Semantic level of interoperability is also affected by the local information retrieval system's functionality and indexing policies. Although the standard provides mechanisms for clearly—if not unambiguously—expressing search requests, retrieval requests, and other IR functional requests, the differences in local systems can jeopardize semantic interoperability. In the example above, the two systems are online library catalogs (i.e., bibliographic databases) populated with records derived from standard MARC records. However, System A allows specific MARC fields to be searched for corporate author names while System B, with the same basic set of records, has chosen not to create indexes or is incapable of creating indexes to support the

access point of corporate author. Thus System B is incapable of doing a search for corporate author even though the Z39.50 server front end to its system can process and understand the query. There is likely a strong relationship of the search capabilities of the underlying IR system and the Z39.50 attributes it supports in its Z39.50 server software. Further, Z39.50 client and server software cannot add functionality to a local IR system that it doesn't have.

As a community, we are beginning to grasp the impact of local systems' functionality, local indexing decisions and policies, normalization practices, etc., on interoperability. These impacts go beyond issues of Z39.50 conformance but part of the interoperability equation can be addressed by Z39.50 profiles.

Z39.50 Profiles: Solutions to Semantic Interoperability

Profiles can be considered auxiliary standards mechanisms. They define a subset of specifications from one or more standards to improve interoperability. The objective of a profile is to detail a set of specifications from options and choices available in a base standard(s) to address specific technical or functional requirements. Implementors' products conforming to a profile have an improved likelihood of interoperability. Two motivations have initiated Z39.50 profiles:

- to prescribe how Z39.50 should be used in a particular application environment (e.g., government information, cultural heritage museums, etc.)
- to solve interoperability problems with existing Z39.50 implementations within a community or across two or more communities (e.g., the library community).

This section discusses how profiles can address semantic interoperability problems in cross-catalog searching.

Between 1999 and 2000, an international effort produced *The Bath Profile: An International Z39.50 Specification for Library Applications and Resource Discovery* [15, 16]. The Bath Profile itself was informed by several previous profiles, but most importantly by the *Z Texas Profile: A Z39.50 Profile for Library Systems Applications in Texas* [17, 18]. These two profiles focused effort on resolving semantic interoperability problems for cross–catalog information retrieval, and they prescribed the specific Z39.50 services required to support various user tasks (e.g., Init, Search, Present, Scan).

In the case of the Bath Profile, it addresses semantic interoperability for searching by defining a core set of 19 searches; requirements for these cross–catalog searches resulted from discussions among librarians. Defining the searches included naming a search, prescribing IR system behavior to process the query, and prescribing the Z39.50 query vocabulary to unambiguously express each defined search. For example, the Profile defines an ***Author Keyword Search with Right Truncation***. The semantics (i.e., prescribed IR system behavior) for that search is: "Searches for complete word beginning with the specified character string in fields that contain the name of a person or entity responsible for a resource." The specification of the query using Z39.50 Attributes is:

- Use Attribute (1) = author (1003)
- Relation Attribute (2) = equal (3)
- Position Attribute (3) = any position in field (3)
- Structure Attribute (4) = word (2)
- Truncation Attribute (5) = right truncation (1)

- Completeness Attribute (6) = incomplete subfield (1).

This combination of attribute types and attribute values expresses this and only this search. Thus, there should not be any ambiguity of what a server is to do when it receives this query. If the Z39.50 server and its database is unable to understand this query or to process it in the way prescribed, it should fail the search and return a diagnostic to the Z39.50 client.

Even though the profiles address the Z39.50 aspect of semantic interoperability, the semantic level is also affected by the indexing policies and search functionality in the local IR system. To address the variations in indexing in different systems, the approach of the Texas Z39.50 Implementors Group (TZIG) is to recommend a common indexing policy to support the searches specified in the Profile. Recommending indexing policies goes beyond the scope of Z39.50 specifications, but to improve semantic interoperability, we have concluded that common indexes populated with data from a core set of MARC fields and subfields is essential.

The library community is quite homogeneous, especially in terms of its catalogs. But the diversity -- in Z39.50 implementations and local information retrieval systems -- is now reducing the ability of users (whether information professionals or end user patrons) to take advantage of the networked environment to discover and retrieve pertinent resources. The experience with the Bath and Z Texas profiles suggest that a new level of standardization and consistency in Z39.50 implementation, information retrieval functionality, and indexing practices is necessary to achieve meaningful networked information retrieval among library catalogs.

Virtual Union Catalogs and Cross-Domain Searching

The final sections of this paper present two applications areas in which Z39.50 is being used currently. These fall generally into the arena of resource discovery since these applications involve the identification of an information resource for retrieval and access.

Virtual Union Catalogs

Although the original model of intersystem communication for Z39.50 focused on a Z39.50 client interacting with a single Z39.50 server, implementors in the 1990s began developing clients that allowed a user to interact with more than one Z39.50 server at a time. This gave the user the capability of formulating a single search that would be executed against two or more separate databases. The Z39.50 client established Z39.50 sessions with one or more servers, sent the query to each of those servers, and retrieved results from each server to present to the user. >From the user's perspective, he/she was simultaneously and transparently searching multiple resources at the same time. As a result, the multiple resources being searched at the same time appeared to the user as a single search against one logical resource.

Librarians saw the potential for this in the context of union catalogs [19]. Why not use the distributed searching capabilities of Z39.50 to create virtual union catalogs by virtue of sending the same query to multiple catalogs simultaneously? Would it be possible to abandon the physical union catalog in favor of a virtual union catalog? Figure 4 illustrates how a Z39.50 client connects to multiple, remote catalogs for search and retrieval. A single search from the user is sent to multiple Z39.50-accessible catalogs and results from each catalog are returned. Depending on the client-side capabilities, the results from each

of the catalogs could be merged into a single result set with duplicate records removed, etc. From the users' perspective, however, the search goes against a logical resource (i.e., the virtual union catalog) rather than against separate catalogs.

Figure 4
Virtual Union Catalog Application

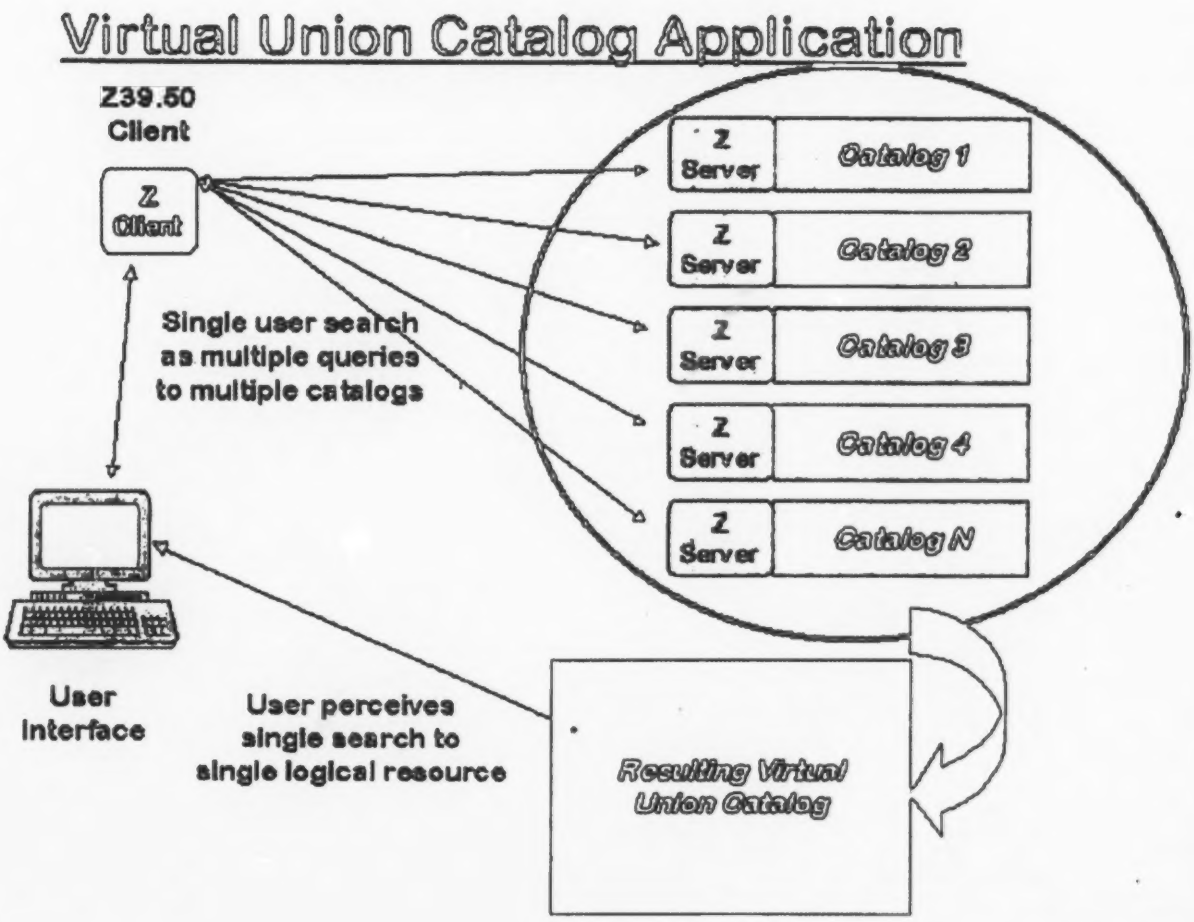

The use of Z39.50 doesn't mean the end of traditional union catalogs. For example, Clifford Lynch suggests that we should see that the single physical union catalog model "complements the emerging distributed search models by offering substantially different functionality, quality, performance, and management characteristics" [20]. To adequately assess the utility of either model, however, studies are needed to evaluate these differences. Coyle provides one of the first systematic looks in comparing a centralized union catalog (i.e., Melvyl) with a virtual union catalog [21].

Performance issues may become paramount considerations. For example, in a virtual union catalog each search will go to each participating catalog. Smaller public libraries participating in such a catalog may be subject to large numbers of virtual union catalog search that could put an adverse load on local computing resources compared to a large academic library participant with a more robust computing and networking infrastructure. Performance issues have yet to be investigated systematically.

BEST COPY AVAILABLE

And we also have to deal with the ever-present semantic interoperability issues in a virtual union catalog model. Unless each participating catalog's Z39.50 server is configured similarly for support of Z39.50 attribute types and values, and each catalog's indexing policies are similar, users may be less satisfied with the results from a virtual union catalog than from a centralized single union catalog database [19]. These semantic interoperability problems, however, are susceptible to the solutions provided by Z39.50 profiles.

Cross Domain Searching

Library catalogs are not the only resources that are Z39.50 accessible. Efforts in the cultural heritage museum, natural history museum, archives, government information, and geospatial communities to implement Z39.50 solutions for networked information retrieval are making a diverse set of information resources available to Z39.50 clients. It may be that when one thinks of the concept resource discovery, this heterogeneous networked information environment is what captures their imagination. Think of a user with a need for information about the artist Van Gogh. Certainly the user might be interested in discovering books about the artist, but he/she might also be interested in discovering manuscript collections, images, museum collections and exhibits, etc. related to Van Gogh. The user might begin with a search of several library catalogs plus one or more museum systems and an archive or other metadata repository to find relevant information. Librarians and library users desire integrated access to distributed resources where those resources may take different forms (e.g., images, books, sound recordings, etc.). As Hammer noted, "The essential power of Z39.50 is that it allows diverse information resources to look and act the same to the individual user" [22]. Is this, then, really the promise of Z39.50 and resource discovery?

Z39.50 can be used to provide effective cross-domain searching of diverse resources including library catalogs, government information, museum systems, and archives. A library's Z39.50 client configured for cross-domain searching could send out queries to Z39.50 accessible museum and archive systems configured to support cross-domain searching. Similarly, a museum curator could use a museum Z39.50 client configured to support cross-domain searching to search the local museum system, one or more other museum systems, one or more library catalogs, and government resources that are Z39.50 accessible and configured to support cross-domain searching. A project conducted by the Consortium for the Computer Interchange of Museum Information (CIMI) demonstrated how cross-domain searching could be done across library catalogs and museum collections [23].

One mechanism to enhance Z39.50 cross-domain searching is to use the Dublin Core Metadata Elements to provide semantic interoperability for expressing search requests and packaging retrieval results. In the virtual union catalog described above, there is a homogeneity to the bibliographic records in each catalog (e.g., most all records have a concept of author, title, etc.; they can be interchanged as MARC records). When one moves outside a single domain, that homogeneity of semantics and data structures is removed. In a museum's collection management system, the person responsible for the intellectual work of a painting is seldom referred to as an author but more likely as artist. Yet there is a level of semantic equivalence between the concepts *author* and *artist*.

The Dublin Core Metadata Elements address semantic interoperability for resource discovery [24]. The elements themselves can be used as the "words" in the Z39.50 query vocabulary (i.e., as Use Attributes in Z39.50 to be able to characterize search terms). The Dublin Core elements become a lens through which a Z39.50 client sees a wide range of diverse resources. Similarly, an information retrieval system can make its resource visible through the Dublin Core elements. For retrieval purpose, a Z39.50 server can package up a retrieval record using the Dublin Core elements as labels for the units of information or fields of the retrieval record. Figure 5 illustrates how cross-domain searching can be enabled through the use of Dublin Core elements.

Figure 5
Cross Domain IR Application

Cross Domain IR Application

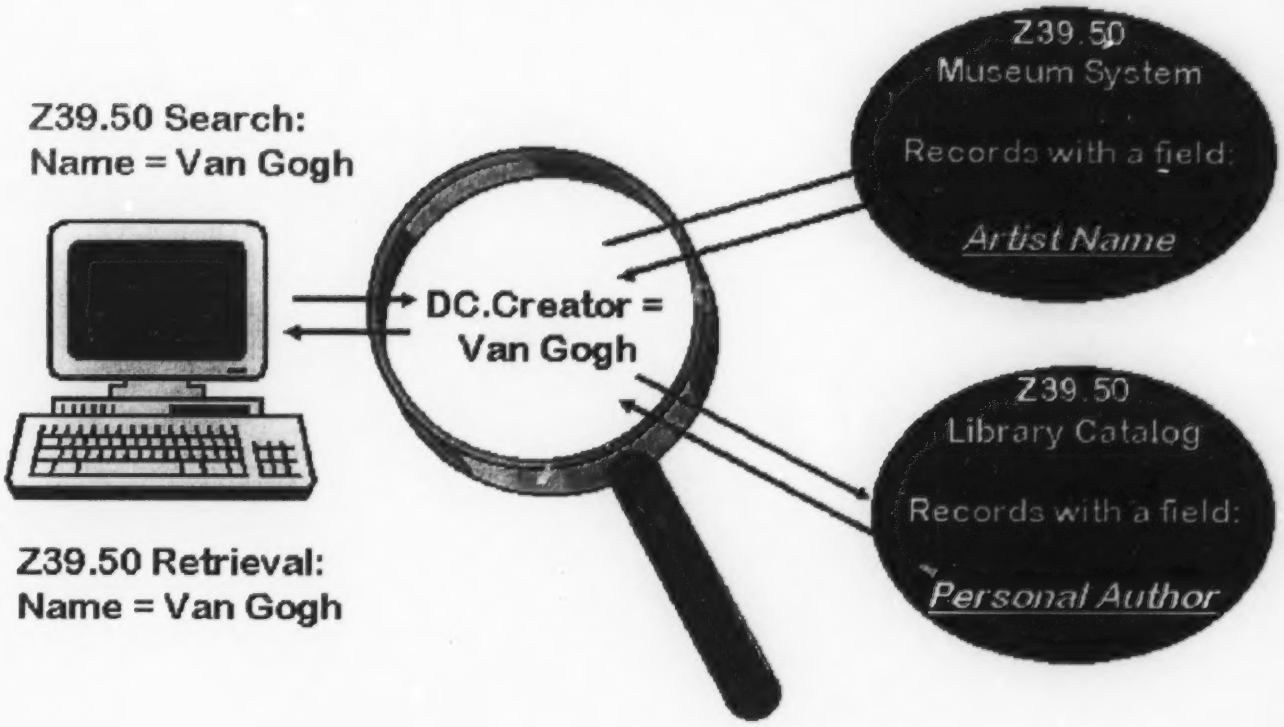

While most of this paper has focused on interoperability issues related to searching, there is an associated set of issues related to retrieval interoperability. In the cross-domain environment, retrieval issues become much more pronounced than in the virtual union catalog. In the latter, retrieval interoperability is achieved through the use of a MARC record syntax for the retrieval record. Most library catalogs can export legitimate MARC records, and these can be passed between the server and client via Z39.50.

Searching across domains, however, offers no such pre-existing standard for a data interchange format. Z39.50 developers addressed this problem in the early 1990s by defining a Generic Record Syntax (GRS) to express arbitrarily structured database records in a standard format for interchange in Z39.50. While this proved to be a viable solution within the Z39.50 community, a more likely solution is the integration of Extensible Markup Language (XML) as a core record syntax for use in Z39.50. Whether GRS or XML, addressing semantic interoperability on the retrieval side is as pressing as the semantic interoperability on the searching side when doing cross-domain searching.

Z39.50's Future in Networked Information Retrieval

The ANSI/NISO Z39.50 protocol for information retrieval is considered by some as an important strategic tool for providing

BEST COPY AVAILABLE

integrated access to distributed networked resources. Others, however, consider it to be an outdated "technology" that should be abandoned. Assessing its utility necessitates a clear statement of the application and functional requirements in which Z39.50 is being considered. Clear functional requirements for an application can then allow us to determine if Z39.50 or some alternative technology is appropriate.

This paper has briefly reviewed the 20+ year history of Z39.50 development, the complexity of information retrieval problems it addresses, and how the goals for its use has changed over time. This standard -- intended to solve problems within a limited community (i.e., libraries) -- now is deployed in a range of other communities to solve the challenges of networked information retrieval. The standard can be viewed as a class of evolutionary standards, and it has evolved to incorporate advances in technologies and technical approaches (e.g., the use of the Internet, integration into the Web environment, and use of new technologies such as the XML).

Where does the perception that Z39.50 represents outdated technology arise? Without some attention to this issue, any discussion of Z39.50's future is clouded. Z39.50's origins in the Open Systems Interconnection (OSI) framework of the 1970s and 1980s have not been forgotten (nor entirely removed from the standard). The power of Z39.50 comes at a cost of complexity. Setting up a web server and full-text indexing search engine is commonplace. How common is it for an operating system to bundle an easy-to-configure Z39.50 server as, for example, Linux does with the Apache web server? Available Z39.50 toolkits may require not only significant C or C++ programming experience but also require familiarity with the less-than-common technical tools such as Abstract Syntax Notional One (ASN.1) and Basic Encoding Rules (BER) to encode the protocol messages for transmission over the wire. A Z39.50 implementor has to address a range of concerns from abstract semantic models to the bits passing over the wire. And, for the most part, there is little off-the-shelf software that can make implementing Z39.50 clients or servers easy to do. Certainly we don't see Z39.50 plug-ins for Netscape and Internet Explorer.

Will Z39.50 be relegated to a backwater of networked information retrieval? It is a standard that addresses important interoperability challenges but does so in a way, perceived as a library way, that may keep it a niche solution rather than as a broader solution to critical problems of networked information retrieval. This paper has argued that major contributions of Z39.50 have been abstract and semantic models for information retrieval. The question is whether and how the Z39.50 community can leverage these contributions while letting go of some of the arcane technical aspects of the protocol that keep it from being widely adopted. At the July 2000 international Z39.50 Implementors Group (ZIG) meeting in Leuven, Belgium, participants agreed to build on the strengths of Z39.50 (the modeling and abstraction) and investigate how other technologies and newer protocols could be used (e.g., SOAP and the emerging XML Protocols).

Z39.50's future in broader networked information retrieval environment is uncertain. The complexity of distributed networked information retrieval is not appreciated until one tries to do it. Information retrieval from a single IR system is not problematic (as is the case with the web search engines). Distributed search across multiple servers with different database systems and different data and semantic structures is problematic. Experience with Z39.50 has identified many aspects of the complexity of distributed search and retrieval. Z39.50 developers and implementors have worked to resolve many interoperability issues, but too often the successes have come slowly and usually not with great fanfare.

The strategy for success being followed by the Bath and Z Texas Profile developers may be considered an incremental strategy. We are trying to rebuild confidence in Z39.50 for a group of users that should not have lost confidence in the first place, namely, librarians. We are not promising that Z39.50 will solve all information retrieval problems. But the profiles offer an opportunity to show how Z39.50 can be used successfully in the original community that developed the standard. Discussing Z39.50's role in resource discovery as compared with web search engines, although attempted in this paper, may be one more tangent from the pragmatic roles for Z39.50:

- as a standard that provides an example of mechanisms for "standardizing shared semantic knowledge" [4]
- as a practical tool in the arsenal of librarians and information professionals in search and retrieval across multiple library catalogs
- as a potential strategic tool for integrating access to selected networked information resources.

Success in these three roles is possible. Demonstrable and effective use of Z39.50 within the library community has not been a

given. We can at least start Z39.50's future by making it work for us in the present.

Notes

1. For a history of Z39.50 development, see Moen, William E. (1998). *The development of ANSI/NISO Z39.50: A case study in standards evolution*. Unpublished doctoral dissertation. Syracuse, NY: School of Information Studies, Syracuse University. Available: <http://www.unt.edu/wmoen/dissertation/DissertationIndex.htm>.
2. The National Information Standards Organization (NISO) resulted from an organizational change in 1984. Prior to that point, the name of the standards organization was the American National Standards Committee Z39. This paper will use NISO to reference the current as well as earlier standards organization.
3. American National Standards Committee Z39. (1979, February 2). *Information form: Recommended ANSC Z39 Standard*. Gaithersburg, MD: American National Standards Committee Z39.
4. Lynch, Clifford A. (1997, April). The Z39.50 information retrieval standard: Part I: A strategic view of its past, present and future. *DLib Magazine*. Available: <http://www.dlib.org/dlib/april97/04lynch.html>.
5. Wake, William. (2000, February). *Analysis objective: Z39.50 search system - object model.* Available: <http://users.vnet.net/wwake/analysis/z3950/z3950.shtml>.
6. Semantic interoperability is viewed as a difficult problem to solve. For example, Clifford Lynch and Hector Garcia-Molina stated: "Deep semantic interoperability is a 'grand challenge' research problem; it is extraordinarily difficult, but of transcendent importance, if digital libraries are to live up to their long-term potential." Lynch, Clifford and Garcia-Molina, Hector. (1995). *Interoperability, scaling, and the digital libraries research agenda: A report on the May 18-19, 1995 IITA digital libraries workshop.* (U.S. Government's Information Infrastructure Technology and Applications (IITA) Working Group, Reston, Virginia). Available: <http://www-diglib.stanford.edu/diglib/pub/reports/iita-dlw/>.
7. Lynch, Clifford A. (1995, October). Networked information resource discovery: An overview of current issues. *IEEE Journal on Selected Areas in Communications, 13*(8): 1505-1522.
8. Arts and Humanities Data Service. (2000). *The AHDS gateway.* Available: <http://ahds.ac.uk:8080/ahds_live/>.
9. For a longer exposition on interoperability in the context of Z39.50, see Moen, William E. (forthcoming). "Assuring interoperability in the networked environment: Standards, evaluation, testbeds, and politics. In McClure, Charles R. and Bertot, John Carlo, Eds. Evaluating networked information services: Techniques, policy, and issues. Silver Spring, MD: American Society for Information Science, Information Today, Inc.
10. Lynch, Clifford. (1993, March) Interoperability: The standards challenge for the 1990s. *Wilson Library Bulletin, 67*(7): 38-42.
11. Payette, Sandra, Blanchi, Christophe, Lagoze, Lagoze, Overly, Edward A. (1999, May). Interoperability for digital objects and repositories. *D-Lib Magazine,* 5(5). Available: <http://www.dlib.org/dlib/may99/payette/05payette.html>.
12. Miller, Paul. (2000, June 21). Interoperability: What is it and why should I want it? *Ariadne* 24. Available: <http://www.ariadne.ac.uk/issue24/interoperability/intro.html>.
13. Moen, William E. (2000). Interoperability for information access: Technical standards and policy considerations. *The Journal of Academic Librarianship, 26*(2): 129-132.
14. Abbas, June, Monika Antonelli, Mark Gilman, Pamiela Hight, Valli Hoski, Jodi Kearns, Teresa Lepchenske, Martha Peet, Mike Pullin, and Amy Stults. (1999). *An Overview of Z39.50, supplemented by a case study of implementing the Zebra server under the Linux operating system.* Denton, TX: School of Library and Information Sciences, University of North Texas.
    Available: <http://www.unt.edu/wmoen/Z3950/GIZMO/contents.htm>.
15. Bath Profile Group. (2000, June). *The bath profile: An international Z39.50 specification for library applications and resource discovery, Release 1.1. An internationally registered profile.* Available: <http://www.ukoln.ac.uk/interop-focus/bath/current/>.
16. For background on the Bath Profile see Lunau, Carrol. (2000, March). *The Bath profile: What*

is it and why should I care? Ottawa, Canada: National Library of Canada. Available: <http://www.nlc-bnc.ca/bath/prof.pdf>.

17. Texas Z39.50 Implementors Group. (1999, April). *Z Texas profile: A Z39.50 profile for library systems applications in Texas, Release 1.0.* Available: <http://www.tsl.state.tx.us/ld/projects/z3950/TZIGProfile99Apr20.htm>.
18. For background on the Z Texas Profile, see Moen, William E. (1998a). *Texas Z: The Texas Z39.50 requirements and specifications project. A discussion paper*. Prepared for the Texas State Library and Archives Commission. Available: <http://www.unt.edu/wmoen/Z3950/TexasZDPAug98.htm>.
19. See for example, Lunau, Carrol. (1998, June). *Virtual Canadian union catalogue pilot project: Final report.* Ottawa: National Library of Canada. Available: <http://www.nlc-bnc.ca/resource/vcuc/vcfinrep.pdf>.
20. Lynch, Clifford A. (1997, Winter). Building the infrastructure of resource sharing: Union catalogs, distributed search, and cross-database linkage. *Library Trends, 45*(3): p. 449.
21. Coyle, Karen. (2000, March). The virtual union catalog: A comparative study. *DLib Magazine, 6*(3). Available: <http://www.dlib.org/dlib/march00/coyle/03coyle.html>.
22. Hammer, Sebastian and Favaro, John. (1996, March). Z39.50 and the world wide web. *DLib Magazine*. Available: <http://www.dlib.org/dlib/march96/briefings/03indexdata.htm>.
23. Moen, William E. (1998, April). Accessing distributed cultural heritage information. *Communications of the ACM, 41*(4): 45-48. See also the CIMI website at: <http://www.cimi.org>.
24. See Dublin Core Metadata Initiative website at: <http://purl.org/DC/>.

Library of Congress
January 23, 2001
Comments: lcweb@loc.gov

Conference Home Page

What's new

Greetings from the Director for Cataloging

Topical discussion groups

NAS study and 2 articles from the LC staff *Gazette*

Conference program

Speakers, commentators, and papers

Conference sponsors

Conference discussion list

Logistical information for conference participants

Conference Organizing Team

Bicentennial Conference on Bibliographic Control for the New Millennium

Confronting the Challenges of Networked Resources and the Web

sponsored by the Library of Congress Cataloging Directorate

William E. Moen, Ph.D.
Assistant Professor
School of Library and Information Sciences
University of North Texas
P.O. Box 311068
Denton, TX 76203-1068

Resource Discovery Using Z39.50: Promise and Reality

About the presenter:

Dr. William Moen teaches courses in the School of Library and Information Sciences, University of North Texas, on the information of organization, metadata and networked information organization and retrieval, and Z39.50 His research program includes the organization of networked resources, interoperability testing; distributed searching and networked information retrieval, metadata and networked information evaluation, user studies related to networked information seeking behavior, and the development and implementation of technical standards. Since the early 1990s, he has been involved in a number of Z39.50 initiatives including the development of the Government Information Locator Service (GILS) and its associated Z39.50 profile, the coordination of a Z39.50 profile for search and retrieval of cultural heritage information for the Consortium for the Computer Interchange of Museum Information (CIMI), and a comprehensive evaluation of the U.S. Federal implementation of GILS. Most recently, Moen facilitated the development of a statewide Z39.50 profile for Texas (the Z Texas Profile) and participated on the Bath Profile,

BEST COPY AVAILABLE

Cataloging Directorate Home Page

Library of Congress Home Page

and International Z39.50 specification for Library Applications and Resource Discovery. He serves as editor of the CIMI, Z Texas, and Bath Profiles. Recent writings include: *Assuring Interoperability in the Networked Environment: Standards, Evaluation, Testbeds, and Politics* (forthcoming) and articles on Z39.50 and related issues published in the *Journal of Academic Librarianship*, *Journal of the American Society for Information Science*, *Texas Library Journal*, and *Communications of the ACM*. Dr. Moen received his Ph.D. from Syracuse University where he wrote his dissertation on the development of the Z39.50 standard.

Full text of paper is available

Summary: The ANSI/NISO Z39.50 protocol for information retrieval is considered by some as an important strategic tool for providing integrated access to distributed networked resources in the while others consider it to be an outdated "technology" that should be abandoned. An understanding of its historical development is critical to evaluate the current perceptions and misperceptions of the roles it is assuming in the networked environment. This paper briefly reviews the 20+ history of Z39.50 development, the complexity of information retrieval problems it addresses, and how the goals for its use has changed over time. In part, the paper shows how this standard was intended to solve problems within a limited community (i.e., libraries) but has now become deployed in other communities to solve the challenges of networked information retrieval. The standard can be viewed as a class of evolutionary standards, and it has evolved to incorporate advances in technologies and technical approaches (e.g., the use of the Internet, integration into the Web environment, and use of new technologies such as the Extensible Markup Language).

The context of Z39.50's goals provides a way to investigate the meaning of resource discovery. Like many terms in the networked environment, resource discovery has many meanings, and the paper attempts to identify the type of resource discovery enabled by Z39.50. Networked resource discovery implies the use of one system to discover resources on one or more separate systems, and such interworking of two systems highlights the key issue of interoperability.

One constant goal of Z39.50 developers was to enable interoperability between diverse systems and diverse resources. The paper describes how Z39.50 enables this interoperability yet details reasons why implementations of the standard have been deficient in achieving this important goal. Recent initiatives have resulted in important national and internationals specifications for using Z39.50 (i.e., profiles) to address underlying interoperability problems, and profiles appear to offer a realistic solution path for seemingly intractable problems in interoperability. The paper describes these profiles and the likely impact they will have on the use of Z39.50 both within libraries and within other communities such as museums. In addition, the paper suggests a framework for analyzing the complexity of interoperability and identifies an approach being developed at the University of North Texas for establishing a rigorous interoperability

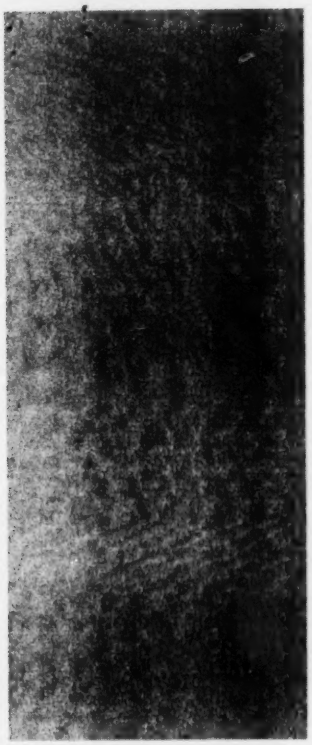

testbed.

The past several years has seen a new uptake of Z39.50, both within the library community for creating virtual union catalogs as well as in other communities to solve networked information retrieval problems and provide services to customers. The paper highlights several of these developments to indicate potential roles for Z39.50 in the networked environment. The paper concludes with an overall assessment of Z39.50 strengths as well as the opportunities and challenges the standard faces in serving as a strategic information retrieval tool for libraries and other communities in the networked environment.

Z39.50 continues to evolve as a comprehensive international standard designed to improve the information retrieval of networked resources in a distributed environment, with examples of numerous "profiles" that have been developed over the last several years. This presentation addresses the perception that the standard lacks the broad Internet community support and the contention that it is too flexible and too large and complex for widespread commercial application. It identifies outstanding problems and looks at how well positioned the standard is to offer a future solution to increasing retrieval problems of networked resources on the Web.

Library of Congress
May 9, 2000
Comments: lcweb@loc.gov